

Fighting Forgery with SPF



Have you noticed that almost all spam comes with a forged envelope sender? SPF gives you a way to detect these forgeries. Anti-spam tools can score messages according to their SPF status. MTAs may also choose to reject messages that fail SPF tests.

How It Works

First a domain owner publishes an SPF record in DNS. The record describes the servers which send mail from that domain – servers which are permitted to use that domain name in the “MAIL FROM” envelope sender return-path in an SMTP transaction. Then receiving MTAs (or downstream tools like SpamAssassin) can look up that record and tell if the sender was legitimate. If that lookup occurs during the SMTP transaction, an MTA can reject forgeries even before DATA is transferred. Think of it as caller ID for email.

History

In 1998 Paul Vixie wrote a paper, “Reputing Mail-From”. In 2003, RMX and DMP, two competing schemes, expanded on his ideas. SPF is a hybrid of both schemes. It has evolved under peer review and will be presented at the Seoul IETF in February 2004 for RFC consideration. Since December, when the design was frozen, many thousands of domains have published records. SPF is catching actual spam today.

Known Issues

Forwarding. SPF breaks Unix `/etc/aliases-style` forwarding because the authentication model cannot accommodate legitimate intermediaries in a decentralized fashion. This is a burden for forwarding services like acm.org and pobox.com, but it is not insurmountable: you can solve it by rewriting the sender address, in a scheme called SRS which is similar to VERP. Pobox.com already does this for customers. SRS patches for MTAs are being developed. Forwarding sites should be able to solve this by simply upgrading their MTA.

Legitimate forgery. SPF won't let you mail through your Earthlink ISP gateway if your return-path is set to your Yahoo address. This problem also affects “mail me this news article” where a third party sends mail on your behalf.

Change. ISPs have to implement SMTP AUTH or POP-before-SMTP so their roaming users can “phone home”. Those roaming users have to put that setting into their MUA. This takes time.

Record Type. An SPF DNS RR type is preferable, but getting new RR types approved can take years. TXT is a decent workaround.

Alternatives. Sender authentication is a good idea, but the designated sender scheme isn't the only way to do it. Embedding credentials in the message itself solves the forwarding problem, though at a higher cost in bandwidth, CPU, and MTA upgrades. The SPF grammar is extensible and will support crypto schemes.

The Big Picture. SPF breaks a number of things. Is the benefit to sender domains (preventing some forgeries) and to receivers (blocking some spam, worms, and viruses that could probably be blocked using other means) worth breaking forwarding and legitimate forgery?

Who's Using SPF?

Domains Publishing since Dec 15 ♦ AOL ♦ AltaVista ♦ DynDNS ♦ LiveJournal ♦ O'Reilly ♦ Oxford.ac.uk ♦ PhilZimmermann.com ♦ Perl.org ♦ w3c.org ♦ *SPF-aware products* ♦ SpamAssassin 2.70 ♦ Sophos PureMessage ♦ Declude JunkMail ♦ MailArmory ♦ MTAs ♦ Postfix plugin ♦ Sendmail milter ♦ Exim ACL ♦ Qmail patch

Adoption Path

At first, a domain can publish a *?all* default meaning *unknown*. This default accommodates users who do not yet send mail through the domain's supported mail gateways, and gives everybody time to adjust and experiment. Eventually, to enjoy the full benefits of SPF protection, the domain should change the default to *-all*, for *fail*.

Elegance

Most anti-spam approaches need to analyze the message data. SPF can operate *before* the SMTP DATA is sent, saving bandwidth and CPU, two major costs felt by ISPs everywhere. SPF strengthens RHSBLs and address verification techniques. SPF is also extensible: future authentication mechanisms can be expressed in the existing grammar.

Why People Use SPF

Most people agree that while it isn't perfect, the pros outweigh the cons. Thousands of domains, not wanting to be hurt by fraud, have published records. Besides, some ISPs have already started doing “pseudo-SPF” to try to detect Yahoo and AOL forgeries. Isn't it better to have a free and open standard that everyone can use? Some critics say the language is complex and hard to parse, but writing plugins for the handful of widely used MTAs is a one-time chore (and has already been done by the opensource community). SPF optimizes for ease of record creation, lowering the adoption barrier for the millions of domains out there. To them, SPF is not complex, but simple!

Reputation and Accreditation

Why are DNSBLs more popular than RHSBLs? Because spammers forge sender addresses. If SPF can get them to expose their true names, they will have to churn through disposable domains which tools can automatically blacklist. Soon, distributed reputation systems will arise that don't give just a binary yes/no answer but can tell you how long a domain has been around and give you a count of total mail volume and number of reported complaints, leaving the final decision up to you. *Bonded-sender.org-style* bulk mailer accreditation will also work well with no need for an IP list.

